

Berkeley HRM Admin guide

Lawrence Berkeley National Laboratory
Scientific Data Management Research Group
Junmin Gu, Arie Shoshani, Alex Sim
April 28, 2004

All the licenses are specified in the DataMover package.

HRM consists of two components; DRM and TRM. DRM manages the disk and TRM manages the access to the MSS. DRM can be run by itself without TRM for its disk management capability.

DRM Admin Guide:

We use *drmExecutable* to denote drmServer.sol or drmServer.linux, depending on your operating system.

Configuration

To use DRM, do "*drmExecutable* *your_rc_file*"

DRM supports all corba arguments, and they can be just provided at the command line. For instance, if you want to run DRM at a specific port number because of the firewall issues, -OApport port_number (e.g. -OApport 6172) to be added as part of the command line. Another useful option is -OAthread_pool int_num, that it increases the number of thread pool depending on your machine spec. (e.g. -OAthread_pool 25)

To exit DRM. press "CTRL-C"

To see all options, try drmServer.sol(or drmServer.linux), and it lists a number of options supported in the *drm** section of configuration file . Here is a detailed description:

CORBA related parameters:

DRM registers itself to a given naming service using a given object name. The address of naming service can be specified in *drm*nameservice*. If not provided, DRM looks for values of *common*NSHost* & *common*NSPort* in the configuration file.

*drm*nameservice*
*drm*objectname* (default: DRMServer)
*drm*threadtype* (default: thread_pool)

Cache management options:

DRM finds the path of its managed cache directory from *drm*cachehome*, and size of from *drm*maxCacheSize*. Usually a file request provides the size of a file. If not, then DRM will assume a default value *drm*defaultFileSize*. If user provided file size is smaller than the actual

size, DRM will increment the space allocation by `drm*fileIncrement`, until file size reaches `drm*maxFileSize`.

`drm*cachehome` (no default. must supply)
`drm*maxCacheSize` (examples: 1GB, 500MB, 7000000KB)
`drm*maxFileSize` (upto maxCacheSize or 2GB)
`drm*fileIncrement` (upto maxFileSize or 1GB)
`drm*defaultFileSize` (upto maxFileSize)
`drm*autoSuspendTime` (default: 250(in seconds))
`drm*enableFileSharing` (default: false)
`drm*cacheFileNamingMode` (ENDFILENAME/SUFFIX_ENDFILENAME default: ENDFILENAME)

Function specific options:

When it is set to true, DRM is serviced as a part of HRM, and will look for TRM location through `common*TRMName` in naming service provided by `common*NSHost` and `common*NSPort`.

`drm*useTRM` (default: false)

Log options:

DRM saves several logs for its reference. A file specified by `drm*logfile` will be used to save the file information in cache. Outputs will be saved in `drm*outputFile`, and detailed events in DRM will be saved in `drm*eventLogFile`. These options accepts a full file path. By default, they will be created in the same directory where the DRM executable is running. Preferably, these log files should not be on the same disk where the DRM managed cache is located.

`drm*logfile` (default: log.cache)
`drm*outputFile` (default: assisigned by DRM)
`drm*eventLogFile` (default: assisigned by DRM)
`drm*appendExistingEventLogs` (default: false)
`drm*showDebugMessages` (default: false)

Processing options:

The parameter `drm*concurrencyLevel` denotes how many file requests can be processed concurrently , while `drm*maxConcurrentFTP` controls how many file transfers can happen concurrently . If `drm*generalService` is false, then DRM will accept all users. Otherwise, only users in the local gridmap will be able to access DRM. If a request is sent to another SRM by DRM for a file request, DRM will send status calls to another SRM until the request is processed. The time between two status calls are decided by `drm*checkpoint`. If a file request is failed but needs to be tried again, DRM would retry after some delay, by `drm*retrypoint`. When a file is transferred to its managed cache, it will be pinned for the user. Each user has their own pinning limit, by `drm*maxSizeConcurrentPinnedFiles` and `drm*maxNumConcurrentPinnedFiles`, whichever number is reached first. To set the length of each pin in seconds, use

`drm*pinDuration`. The `drm*prefetchingLimit` determines how much space a user is entitled to use in DRM for their files (pinned and unpinned.)

`drm*maxConcurrentFTP` (no default)
`drm*concurrencyLevel` (default: 5)
`drm*checkpoint` (default: 200(in seconds))
`drm*retrypoint` (default: 300(in seconds))

`drm*generalService` (default: true)

`drm*prefetchingLimit` (default: unlimited)

`drm*maxSizeConcurrentPinnedFiles`
`drm*maxNumConcurrentPinnedFiles` (default: 1)
`drm*pinDuration` (default: 50(in seconds))
`drm*enableFileSharing` (default: false)

GridFTP options:

DRM uses GridFTP to transfer files. `drm*localGSIFTPhostname` contains the GridFTP server hostname that DRM will use when notifying the users about the transferable URL. `drm*FTPBufferSize`, `drm*NumFTPStreams`, and `drm*FTPBlockSize` are GridFTP options that DRM will use when transferring files.

`drm*localGSIFTPhostname`
`drm*localGSIFTPhostport` (default: 2811)
`drm*FTPBufferSize` (for globus. default: 0)
`drm*NumFTPStreams` (for globus. default: 0)
`drm*FTPBlockSize` (for globus. default: 0)

Trouble shooting:

If CORBA calls get stuck, try "hrm-ping.sol –conf hrm.rc".

TRM Admin Guide:

TRM manages the access to the Mass Storage System. So far, it manages to HPSS and NCAR-MSS.

We use *trm* to denote trm.sol or trm.linux, depending on your operating system.

-help	print this message
-conf path/file	path for configuration file (default ./hrm.rc)
-mss location	MSS location [LBNL, BNL, ORNL, NCAR, SCP] (default LBNL)
-msss type	MSS security type [GSI, ENCRYPT, KERBEROS, SSH, PLAIN] (default GSI)
-trm name	TRM's server object name (default HPSSResourceManager)
-host hostname	host name for the TRM (default from cfg file)
-port int	port number for TRM (default 5533)
-ref path/file	Reference file path for TRM (default /tmp/trm.ref)
-hsipath	enable HSI and its path (default from cfg file)
-diskpath path	data directory to bring in files (default from cfg file)
-pfiles int	set the number of files to get per tape (default 0)
-debug	enable debugging (debug info displayed) (default true)
-quiet	quiet/silence (no debug info) (default false)
-silence	quiet/silence (no debug info) (default false)
-log path/file	enable logging to a file (default true and /tmp/out.trm.log)
-logdetail	enable detailed logging (default false)
-recovery path/file	enable recovery and path (default false and /tmp/trm.recovery.info)
-ssh path	path to ssh public key for TRM (default \$HOME/.ssh/identity.pub)
-pftppath	path to pftp (default pftp assuming it's on path)
-pftplimit int	maximum number of pftp allowed (default 5)
-passcheck	bypass checking for HPSS/HSI login (default false)
-defaultlogin	set MSS default login and passwd to be prompted (default false)
-msshost hostnames	host names and port number (if any) for HPSS (default from cfg file)
-hpsshsihost hostnames	host names and port number (if any) for HPSS HIS (default from cfg file)
-protocols names	supported file transfer protocols in TRM (delimited by ; if multiple) (default from cfg file)
-numeric	using numeric IP address for hostnames (default false)
-v	show version

All of these options can be defined in the configuration file as well. (Note: Configuration file can be shared with DRM.)

To specify HSI (HPSS access command) path:

trm*HSI=path_to_hsi

Note: For some MSSes, HSI option will not be used and needed. If HSI is not provided, this command will not be used.

To specify MSS hostname for HSI:

trm*HPSSHSIHostName=MSS_HSI_host_name

Note: If *HSI is provided but *HPSSHSIHostName is not, TRM will use MSSHostName instead to access MSS with HSI.

To specify PFTP_Client (HPSS access command) path:

trm*PFTP=path_to_ftp

Note: For some MSSes, PFTP option will not be used and needed.

To specify the local disk path to retrieve files, and it needs to be the same as DRM*cachehome:

trm*DiskPath=DRM_cache_path

To specify the MSS hostname:

trm*MSSHostName=hostname_for_MSS_access

To specify the maximum number of concurrent MSS transaction:

trm*MSSMaximumAllowed=integer

In case MSS does not respond, MSSWaitTime is the re-trial time in seconds:

trm*MSSWaitTime=120

In case MSS does not respond, how many times does TRM re-try the connection:

trm*MSSRetrial=30

To specify the event log file:

trm*TRMLLogFile=/path_to_log/out.trm.log

trm*EnableLogging=true

To have detailed log files, including all MSS transaction logs:

trm*EnableDetailedLog=true

To have recovery option in case TRM crashes:

trm*EnableRecovery=false

trm*RecoveryPath=/path_to_existing_recovery_directory

NCAR-MSS specific command paths for msls and msrccp:

trm*MSLS=path_to_msls

trm*MSRCP=path_to_msrrcp

Sample Configuration FILE

Note: Blank lines and a line with starting # character will be ignored.

Notes on prefixes:

common* : will be read by all components

drm* : will be read only by the DRM

trm* : will be read only by the TRM

hrm* : will be read only by the DataMover client programs.

Sample hrm.rc

```
#####
common*NSHost=srm.lbl.gov
common*NSPort=6171
common*TRMName=HPSSResourceManagerASIM
common*TRMRefFile=/srm/data2/srm/log/trm.ref
common*VerboseConfig=true
#####
drm*objectname=DRMServerOnHRMASIM
drm*cachehome=/home/dm/srm/data2/srm/data
drm*maxCacheSize=10GB
drm*maxConcurrentFTP=3
drm*concurrencyLevel=25
drm*localGSIFTPhostname=srm.lbl.gov
drm*logfile=/home/dm/srm/data2/srm/log/log.cache.hrm
drm*useTRM=true
drm*NumFTPStreams=2
drm*FTPBufferSize=1000000
drm*maxSizeConcurrentPinnedFiles=3GB
drm*maxNumConcurrentPinnedFiles=30
drm*pinDuration=300
drm*outputFile=/home/dm/srm/data2/srm/log
# to show out.drm.log output, not the event logs
drm*showDebugMessages=true
drm*eventLogFile=/home/dm/srm/data2/srm/log/event.drm.log
drm*appendExistingEventLogs=true
# To retain the original file name in the cache
drm*attachSourceFileNameInCache=true
# To set prefetching limit (default=unlimit=0)
drm*prefetchingLimit=500MB
# To enable file sharing
# When it is false (by default), it is good for file replications, but not for client support.
drm*enableFileSharing=true
```

```
drm*cacheFileNamingMode=SUFFIX_ENDFILENAME
#####
# HSI is needed as full path
trm*HSI=/home/software4/hsi/hsi
trm*PFTP=/home/software/pftp-gsi/pftp_client_gsi
trm*DiskPath=/srm/data2/srm/data
trm*MSSHostName=garchive.nersc.gov
trm*MSSMaximumAllowed=3
trm*MSSWaitTime=120
trm*MSSRetrial=30
trm*TRMLLogFile=/srm/data2/srm/log/out.trm.log
trm*EnableLogging=true
trm*EnableDetailedLog=true
trm*EnableRecovery=false
trm*RecoveryPath=/srm/data2/srm/log/trm.recovery
# The followings are NCAR specific
trm*MSLS=/usr/local/bin/msls
trm*MSRCP=/usr/local/bin/msrcp
#####
```

DRAFT